arXiv:1703.00632v1 [cs.GT] 2 Mar 2017

A Dominant Strategy Truthful, Deterministic Multi-Armed Bandit Mechanism with Logarithmic Regret

Divya Padmanabhan*, Satyanath Bhat*, Prabuchandran K.J.†, Shirish Shevade*, and Y. Narahari*

* Indian Institute of Science, Bangalore
† IBM IRL, Bangalore

Abstract. Stochastic multi-armed bandit (MAB) mechanisms are widely used in sponsored search auctions, crowdsourcing, online procurement, etc. Existing stochastic MAB mechanisms with a deterministic payment rule, proposed in the literature, necessarily suffer a regret of $\Omega(T^{2/3})$, where T is the number of time steps. This happens because the existing mechanisms consider the worst case scenario where the means of the agents' stochastic rewards are separated by a very small amount that depends on T. We make, and, exploit the crucial observation that in most scenarios, the separation between the agents' rewards is rarely a function of T. Moreover, in the case that the rewards of the arms are arbitrarily close, the regret contributed by such sub-optimal arms is minimal. Our idea is to allow the center to indicate the resolution, Δ, with which the agents must be distinguished. This immediately leads us to introduce the notion of Δ-Regret. Using sponsored search auctions as a concrete example (the same idea applies for other applications as well), we propose a dominant strategy incentive compatible (DSIC) and individually rational (IR), deterministic MAB mechanism, based on ideas from the Upper Confidence Bound (UCB) family of MAB algorithms. Remarkably, the proposed mechanism Δ-UCB achieves a Δ-regret of $O(\log T)$ for the case of sponsored search auctions. We first establish the results for single slot sponsored search auctions and then non-trivially extend the results to the case where multiple slots are to be allocated.

1 Introduction

Multi-armed bandit (MAB) algorithms [8] are now widely used to model and solve problems where decisions are required to be made sequentially at every time step and there is an *exploration* - *exploitation* dilemma. This dilemma is the tradeoff that the planner faces in deciding whether to explore arms that may yield higher rewards in the future or exploit the arms that have already yielded high rewards in the past. If the rewards are generated from fixed distributions with unknown parameters, the setting goes by the name stochastic MAB [8]. Popular algorithms in the stochastic MAB setting include Upper Confidence Bound (UCB) based algorithms [2] and Thompson Sampling [1] based

algorithms. These algorithms incur $O(\log T)$ regret where T is the total number of rounds or time steps.

When the arms are controlled by *strategic* agents, we need to tackle additional challenges. Mechanism design has been applied in this context, leading to stochastic MAB mechanisms. An immediate example is sponsored search auctions (SSA). In SSA, there are several advertisers who wish to display their ads along with the search results generated in response to a query from an internet user. In the standard model, an advertiser has only one ad to display. We use the terms agent, ad, and advertiser interchangeably. There are two components that are of interest to the planner or the search engine, (1) *stochastic component*: click through rate (CTR) of the ads or the probability that a displayed ad receives a click (2) *strategic component*: valuation of the agent for every click that the agent's ad receives. The search engine would seek to allocate a slot to an ad which has the maximum social welfare (product of click through rate and valuation). However neither the CTRs nor the valuations of the agents are known. This calls for a learning algorithm to learn the stochastic component (click through rate) as well as a mechanism to elicit the strategic component (valuation). This problem could become much harder as the agents may manipulate the learning process [4, 13] to gain higher utilities.

For single slot SSA, it is known that any deterministic MAB mechanism (that is, a MAB mechanism with a deterministic allocation and payment rule) suffers a regret of $\Omega(T^{2/3})$ [4]. Furthermore, there exists a deterministic MAB mechanism with regret matching the theoretical lower bound [4] and also satisfies the strongest notion of truthfulness (a posteriori to the clicks). When a more relaxed notion of truthfulness is targeted (truthfulness in expectation of the clicks), the regret guarantee improves to $O(T^{1/2})$ [3]. The regret can be further improved when randomized mechanisms are used and in fact the regret in this space is $O(\log T)$ [3]. However, the high variance that is inevitable to the payments in randomized mechanisms is a serious deterrent to the use of randomized mechanisms.

We observe that the characterization provided by Babaioff et al. [4] targets the worst case scenario. In particular, in the lower bound proof of $\Omega(T^{2/3})$, they consider an example scenario where the separation, Δ, between the expected rewards of the arms is a function of T. We note that when a similar example ($\Delta = T^{-1}$) is used with the popular UCB algorithm [2], linear regret is obtained, even in the non-strategic case. Hence, a dependence of Δ on T is severely restrictive for the case when the rewards are stochastic, even when the arms are non-strategic. We make the observation that Δ is in most situations independent of T and the planner is knowledgeable enough to specify an appropriate value of Δ. This motivates our main idea in this paper, which is to provide the planner an option to specify a parameter Δ, which is the tolerance or distinguishing level for sub-optimal arms. The understanding is that any arm that is within Δ from the best arm will not cause any additional regret to the planner. This notion of Δ tolerance will require an appropriate definition of regret, which we call Δ-regret. We propose an exploration separated mechanism based on UCB, which achieves

a Δ-regret of $O(\log T)$. This mechanism can be readily applied in several settings such as SSA, crowdsourcing, and online procurement. For the rest of the paper, however, we use SSA as a running example.

Contributions

Our contributions are the following:
(1) We make the crucial observation that in most MAB scenarios, the separation between the agents' rewards is rarely a function of T (the number of time steps). Moreover, in the case that the rewards of the arms are arbitrarily close, the regret contributed by such sub-optimal arms is negligible. We exploit this observation to allow the center to specify the resolution, Δ, with which the agents must be distinguished. We introduce the notion of Δ-Regret to formalize this regret.
(2) Using sponsored search auctions as a concrete example, we propose a dominant strategy incentive compatible (DSIC) and individually rational (IR) MAB mechanism with a deterministic allocation and payment rule, based on ideas from the UCB family of MAB algorithms. The proposed mechanism Δ-UCB achieves a Δ-regret of $O(\log T)$ for the case of single slot sponsored search auctions. The truthfulness achieved by Δ-UCB is a posteriori to the click realizations and is the strongest form of truthfulness.
(3) We non-trivially extend the above results to the case where multiple slots are to be allocated. Here again, our mechanism is DSIC, IR, and achieves a Δ-regret that is $O(\log T)$.

We emphasize that our results are generic to stochastic MAB mechanisms and can be applied to other popular applications such as crowdsourcing and online procurement.

2 Related Work

In the area of MAB mechanisms, a lot of work has been done in sponsored search auctions. Babaioff et al. [4] provide a characterization of truthful MAB mechanisms, wherein the objective is to maximize social welfare. They introduce the notion of influential rounds. The influential rounds are the rounds where the parameters of reward distributions (CTRs) are learnt. One of the characterizations of truthful deterministic mechanisms is that the allocation must be exploration separated, that is, in such influential rounds, the allocation must not depend on the bids of the agents. The allocation is also required to be point wise monotone. One of the main results of their paper is that any truthful, deterministic MAB mechanism incurs a regret of $\Omega(T^{2/3})$. They also provide a mechanism which incurs a matching upper bound regret of $O(T^{2/3})$. Devanur et. al. [9] concurrently provide similar bounds on the regret when the objective is revenue maximization rather than social welfare maximization.

All the above results pertain to the setting of single slot auctions where there is a single slot for which the agents compete. In the generalization of this setting multiple slots are reserved for ads. This setting is more challenging as every

slot is not identical and some slots are more prominent than the others. MAB mechanisms have also been extended to the multiple slot setting [11] in line with the characterization in [4]. Hence, a similar regret of $O(T^{2/3})$ on the social welfare has been attained here as well.

MAB mechanisms have also been proposed in the context of crowdsourcing [7]. Some of these mechanisms incur a regret of $O(\log T)$. This is rendered possible due to the specific nature of the problem in hand. In particular, Bhat et al. [5] look at divisible tasks. Jain et al. [14] look at deterministic mechanisms where a block of tasks is allocated to each agent and provide a weaker notion of truthfulness.

The lower bound of both of social welfare regret as well as regret in the revenue of $\Omega(T^{2/3})$ have influenced subsequent research to follow similar assumptions and thereby obtain a similar regret. However, we show in this work that it is indeed possible to design a deterministic mechanism which attains logarithmic regret and is also truthful in the dominant strategy incentive compatible (DSIC) [15] sense. DSIC, of course, is the most preferred form of truthfulness. This work opens up the possibility for a planner to move away from the worst case scenario to a realistic scenario.

3 The Model: Single Slot SSA

We now describe our SSA setting. Let K be the number of agents or arms. We denote the set of arms by $[K]$. Each of the K arms, when pulled, give rewards from distributions with unknown parameters. We assume here, that the form of the distributions are known but the parameters of the distribution are unknown. In SSA, the rewards of the arms correspond to clicks. The clicks for the advertisements are assumed to be generated from Bernoulli distributions with parameters $\mu_1, \mu_2, \ldots, \mu_K$ where μ_i is the CTR or probability that advertisement i receives a click once observed. The means $\mu_1, \ldots, \mu_K$ are unknown.

A click realization ρ represents the click information of every agent at all rounds, that is, $\rho_i(t) = 1$ if agent i received a click in round t. In a round t, only the click information of the allocated agent is revealed after the completion of the round. Click information of all other unallocated agents is never known to the planner.

The agents also have their valuations for each click they receive. We work in the 'pay per click' setting where the agent pays the search engine for each click received. Let the true valuation of agent i be v_i for a click. v_i is a private type of agent i and is never known to the learner. However the agent is asked to bid his valuation. Let the bid of agent i be b_i. We denote by a vector $b = (b_1, \ldots, b_K)$ the bid profile of all the agents. The central planner wants to ensure that the agents bid their true valuations, that is b_i must be equal to v_i. Assume that there is a single slot which must be allocated to one of the K agents. We denote by W_i the social welfare when agent i is allocated a slot, that is, $W_i = \mu_i v_i$. The social welfare represents the expected valuation of agent i per click. If the CTRs of the agents as well as their valuations were known, the planner would have selected

the arm with the maximum social welfare, that is, $\mu_i v_i$. However neither μ_i nor v_i is known to the planner. Assume v_{max} is the maximum valuation that any agent can have and is common knowledge.

A mechanism $\mathcal{M} = \langle \mathcal{A}, P \rangle$ is a tuple containing an allocation rule $\mathcal{A}$ and a payment rule P. At every time step or round t, the allocation rule acts on a bid profile b of the agents as well as click realization ρ and allocates the slot to one of the K agents, say i. Then $\mathcal{A}(b, \rho, t) = i$. Alternatively we denote the indicator variable $\mathcal{A}_i(b, \rho, t) = \mathbb{1}[\mathcal{A}(b, \rho, t) = i]$. The payment rule $P^t = (P_1^t, P_2^t, \ldots, P_K^t)$, where $P_i^t(b, \rho)$ is the payment to be made by agent i at time t upon receiving a click, when the bids are b and for click realization ρ. As stated earlier $\rho_i(t)$ of the allocated agent alone is observed. Also note that the allocation as well as payments in each round t only depends on the click histories till that round.

Symbol	**Description**
K, $[K]$	No. of agents and agent set
μ_i	CTR of agent i
v_i	Valuation of agent i for each click
W_i	Social welfare when agent i is allocated
$\rho_i(t)$	Click realization of agent i at time t
v_{max}	Maximum valuation over all agents
b_i	Bid of agent i
b	Bid profile of all agents
b_{-i}	Bid profile of all agents except agent i
$N_{i,t}$	No. of times agent i has been selected till time t
$\mathcal{A}(b, \rho, t)$	Allocation at time t for bid profile b and click realization ρ
i_*	Agent with maximum social welfare, ideally must be allocated at every time step
W_*	Social welfare when agent i_* is allocated
Δ	Input parameter by center to indicate the level at which the agents must be distinguished
S_Δ	Set of agents whose social welfare is less than Δ away from i_*. These agents do not contribute to Δ-regret.
$\widehat{\mu}_{i,t}^+$	UCB index corresponding to μ_i at time t
$\widehat{\mu}_{i,t}^-$	LCB index corresponding to μ_i at time t
$\widehat{\mu}_{i,t}$	Empirical CTR of agent i estimated from samples up to time t
P_i^t	Payment charged to agent i if he is allocated a slot at time t and he gets a click

Table 1. Notations for the single slot SSA setting

Let i_* be the arm with the largest social welfare, that is, $i_* = \arg\max_{i \in [K]} W_i$. We denote the corresponding social welfare as $W_* = \max_{i \in [K]} \mu_i v_i$. We denote by I_t the agent chosen at time t as a shorthand for $\mathcal{A}(b, \rho, t)$. For any given $\Delta > 0$, define the set $S_\Delta = \{i \in [K] : W_* - W_i < \Delta\}$. S_Δ denotes the set of all agents separated from the best arm i_* with a social welfare less than Δ. These arms are therefore indistinguishable for the center and they contribute zero to

the regret. Note that Δ is a parameter that the center fixes based on the amount in dollars he is willing to tradeoff for choosing sub-optimal arms, given he has only a fixed time horizon T to his disposal. To capture this revised and more practical notion of regret, we introduce the metric Δ-regret. Formally,

$$\Delta\text{-regret} = \sum_{t=1}^{T} (W_* - W_{I_t}) \mathbb{1}\left[I_t \in [K] \setminus S_\Delta\right] \tag{1}$$

The center suffers a regret only when an agent with a social welfare greater than Δ away from W_* is chosen. Δ-regret captures this loss. The goal of our mechanism is to select agents at every round t to minimize the Δ-regret.

4 Our Mechanism: Δ-UCB

We are now ready to describe our mechanism Δ-UCB. The idea in Δ-UCB is to explore all the arms in a round-robin fashion for a fixed number of rounds. The number of exploration rounds is fixed based on the desired Δ, specified by the planner. At the end of exploration, with high probability, we are guaranteed that the arms not in S_Δ are well separated from the best arm i_* with respect to their social welfare estimates. In the exploration rounds, agents need not pay and these rounds are free.

Further on, for all the remaining rounds, the best arm as per the UCB estimate of social welfare is chosen. However in the exploitation rounds, the chosen agent pays an amount for each click he receives. The amount to be paid by the agent is fixed based on the well known Vickrey Clark Grove (VCG) scheme [16]. Note that no learning place in these rounds and the UCB, LCB indices don't change thereafter. We present our mechanism in Algorithm 1.

5 Properties of Δ-UCB

Next we discuss the properties satisfied by Δ-UCB regarding truthfulness and regret. Before that, we state a few useful definitions which will help in understanding the notion of truthfulness.

At any time step, every agent obtains some utility by participating in the mechanism. This utility is a function of his bid, valuation, bids of other agents and his click realization. Let Θ_i denote the space of bids of agent i. $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_K)$ is the bid profile containing bids of all agents except agent i. Let Θ_{-i} denote the space of bids of all agents other than agent i. Therefore $\Theta_{-i} = \Theta_1 \times \ldots, \times \Theta_{i-1} \times \Theta_{i+1} \times \ldots \times \Theta_K$. We denote by $u_i(b_i, b_{-i}, \rho, t; v_i)$ the utility to agent i at time t when his bid is b_i, his valuation is v_i, the bid profile of the remaining agents is b_{-i} and the click realization is ρ. All agents are assumed to be rational and are interested in maximizing their own utilities.

In our setting the utility to an agent i is computed as,

$$u_i(b_i, b_{-i}, \rho, t; v_i) = (v_i - P_i^t(b, \rho)) \mathcal{A}_i(b_i, b_{-i}, \rho, t) \rho_i(t) \tag{2}$$

Algorithm 1 Δ-UCB Mechanism

Input:

T: Time horizon, K: number of agents

Δ : parameter fixed by the center

v_{max} : Maximum valuation of the agents

Elicit bids $b = (b_1, b_2, \ldots, b_K)$ from all the agents

Initialize $\widehat{\mu}_{i,0} = 0, N_{i,0} = 0 \; \forall i \in [K]$

$u = 8Kv_{max}^2 \log T/\Delta^2$

for $t = 1, \ldots, u$ **do** ▷ Exploration rounds

  $I_t = ((t-1) \mod K) + 1$ ▷ Round-robin exploration

  $N_{I_t,t} = N_{I_t,t-1} + 1$

  $\mathcal{A}(b, \rho, t) = I_t$ ▷ Allocate slot to agent I_t and observe $\rho_{I_t}(t)$

  $\widehat{\mu}_{I_t,t} = (\widehat{\mu}_{I_t,t-1} N_{I_t,t-1} + \rho_{I_t}(t))/N_{I_t,t}$

  $\epsilon_{I_t,t} = \sqrt{2 \log T / N_{I_t,t}}$

  $\widehat{\mu}^+_{I_t,t} = \widehat{\mu}_{I_t,t} + \epsilon_{I_t,t}$

  $\widehat{\mu}^-_{I_t,t} = \widehat{\mu}_{I_t,t} - \epsilon_{I_t,t}$

  $\widehat{\mu}^+_{i,t} = \widehat{\mu}^+_{i,t-1} \; \forall i \in [K] \setminus \{I_t\}$

  $\widehat{\mu}^-_{i,t} = \widehat{\mu}^-_{i,t-1} \; \forall i \in [K] \setminus \{I_t\}$

  $P_i^t(b, \rho) = 0 \; \forall i \in [K]$ ▷ Free rounds

end for

$\hat{i}_* = \arg\max_{i \in [K]} \widehat{\mu}^+_{i,u} b_i$

$j = \arg\max_{i \in [K] \setminus \{\hat{i}_*\}} \widehat{\mu}^+_{i,u} b_i$

$P = \widehat{\mu}^+_{j,u} b_j / \widehat{\mu}^+_{\hat{i}_*,u}$

for $t = u+1, \ldots, T$ **do** ▷ Exploitation rounds

  $\mathcal{A}(b, \rho, t) = \hat{i}_*$

  $P^t_{\hat{i}_*}(b, \rho) = P \times \rho_{\hat{i}_*}(t)$ ▷ Agent pays only for a click

  $P_i^t(b, \rho) = 0 \; \forall i \in [K] \setminus \{\hat{i}_*\}$

  $\widehat{\mu}^+_{i,t} = \widehat{\mu}^+_{i,u}$, $\widehat{\mu}^-_{i,t} = \widehat{\mu}^-_{i,u} \; \forall i \in [K]$ ▷ No more learning

end for

The idea behind the computation of the utility is as follows. If an agent i does not receive an allocation (that is, $\mathcal{A}_i(b_i, b_{-i}, \rho, t) = 0$), his utility is also zero. He gets a non-zero utility only if he receives an allocation. If he receives an allocation and also a click ($\rho_i(t) = 1$), then his utility is the difference between his valuation for the click and the amount he has to pay to the search engine ($v_i - P_i^t(b, \rho)$). If he does not receive a click ($\rho_i(t) = 0$), his utility is zero.

Definition 1. *Dominant Strategy Incentive Compatible (DSIC) [4]: A mechanism $M = \langle \mathcal{A}, P \rangle$ is said to be dominant strategy incentive compatible if $\forall i \in [K], \forall b_i \in \Theta_i, \; \forall b_{-i} \in \Theta_{-i}, \forall \rho, \forall t, u_i(v_i, b_{-i}, \rho, t; v_i) \geq u_i(b_i, b_{-i}, \rho, t; v_i)$.*

Note that in the above definition, the truthfulness is demanded a posteriori to even the click realization [11]. Hence it is the strongest notion of truthfulness. Examples for weaker forms of truthfulness include those which take expectation over click realizations.

Definition 2. *Individually Rational (IR): A mechanism $M = \langle \mathcal{A}, P \rangle$ is said to be individually rational if $\forall i \in [K], \; \forall b_{-i} \in \Theta_{-i}, \forall \rho, \forall t, u_i(v_i, b_{-i}, \rho, t; v_i) \geq 0$.*

Theorem 1. *Δ-UCB mechanism is dominant strategy incentive compatible (DSIC) and individually rational (IR).*

Proof. We analyze the scenarios where an agent i bids his true valuation and receives an allocation and also when he does not. We show that in both these scenarios, bidding his true valuation v_i is indeed a best response strategy. We only need to consider the exploitation rounds because in the exploration rounds, every agent is allocated a fixed number of rounds independent of his bids and these rounds are also free for agents.

Case 1: $\mathcal{A}_i(v_i, b_{-i}, \rho, t) = 1$

This implies that when the agent bids his true valuation, he gets an allocation. Therefore $\widehat{\mu}^+_{i,t} v_i > \widehat{\mu}^+_{l,t} b_l$ for all the other agents l. In particular, let agent j be such that $j = \arg\max_{l \in [K] \setminus \{i\}} \widehat{\mu}^+_{l,t} b_l$. The amount to be paid by agent i is $P^t_i(v_i, b_{-i}, \rho) = \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t}$. If he receives a click then $u_i(v_i, b_{-i}, \rho, t; v_i) = v_i - \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t} > 0$.

Overbid: If agent i bids a value $b_i > v_i$, he continues to receive an allocation and his payment is still the same,
$P^t_i(b_i, b_{-i}, \rho) = \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t}$. Therefore his utility continues to be $u_i(b_i, b_{-i}, \rho, t; v_i) = v_i - \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t} = u_i(v_i, b_{-i}, \rho, t; v_i)$. Therefore he does not benefit from an overbid.

Underbid: Suppose agent i bids a value $b_i < v_i$.

Case a: If b_i is such that $\widehat{\mu}^+_{i,t} b_i < \widehat{\mu}^+_{j,t} b_j$, the he fails to get an allocation as $\mathcal{A}(b_i, b_{-i}, \rho, t) = j \neq i$. Then the utility to agent i is $u_i(b_i, b_{-i}, \rho, t; v_i) = 0 < u_i(v_i, b_{-i}, \rho, t; v_i)$. Therefore he clearly loses his utility by such an underbid.

Case b: Suppose b_i is such that $\widehat{\mu}^+_{i,t} v_i > \widehat{\mu}^+_{i,t} b_i > \widehat{\mu}^+_{j,t} b_j$. That is agent i bids in such a way that he wins the allocation even with an underbid. Then, if he gets a click, the amount he must pay to the center is $P^t_i(b_i, b_{-i}, \rho) = \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t}$. Therefore his utility $u_i(b_i, b_{-i}, \rho, t; v_i) = v_i - \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t} = u_i(v_i, b_{-i}, \rho, t; v_i)$. He obtains the same utility as a truthful bid and there is no benefit from such an underbid.

Case 2: $\mathcal{A}_i(v_i, b_{-i}, \rho, t) = 0$

This implies that when the agent bids his true valuation, he does not get an allocation. Suppose agent j wins the allocation. $\mathcal{A}(v_i, b_{-i}, \rho, t) = j$ and $\widehat{\mu}^+_{i,t} v_i < \widehat{\mu}^+_{j,t} b_j$.

Truthful bid: Since agent i does not win an allocation with a truthful bid, his utility $u_i(v_i, b_{-i}, \rho, t; v_i) = 0$

Overbid: Suppose agent i bids in such a way that $b_i > v_i$. We have two sub-cases here.

Case a: If b_i is such that $\widehat{\mu}^+_{i,t} v_i < \widehat{\mu}^+_{j,t} b_j < \widehat{\mu}^+_{i,t} b_i$, then agent i wins the allocation. So, $\mathcal{A}_i(b_i, b_{-i}, \rho, t) = 1$. If he gets a click, he now has to make a payment $P^t_i(b_i, b_{-i}, \rho) = \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t}$. Now his utility $u_i(b_i, b_{-i}, \rho, t; v_i) = v_i - \widehat{\mu}^+_{j,t} b_j / \widehat{\mu}^+_{i,t} < 0$. And in particular $u_i(b_i, b_{-i}, \rho, t; v_i) < u_i(v_i, b_{-i}, \rho, t; v_i) = 0$. Therefore, such an overbid is clearly disadvantageous compared to a truthful bid.

Case b: Suppose $\widehat{\mu}^{+}_{i,t} v_i < \widehat{\mu}^{+}_{i,t} b_i < \widehat{\mu}^{+}_{j,t} b_j$. The overbid by agent i is not sufficient to make him win the allocation and agent j wins the allocation, $\mathcal{A}(b_i, b_{-i}, \rho, t) = j$. The utility of agent i, $u_i(b_i, b_{-i}, \rho, t; v_i) = 0 = u_i(v_i, b_{-i}, \rho, t; v_i)$. Therefore there is no advantage for agent i by this case of overbid.
Underbid: If agent i bids in such a way that $b_i < v_i$, he continues to lose the allocation and therefore his utility,
$u_i(b_i, b_{-i}, \rho, t; v_i) = 0 = u_i(v_i, b_{-i}, \rho, t; v_i)$. Since, the utility by an underbid remains the same as a truthful bid, there is clearly no advantage in underbidding.

All the above cases show that our mechanism is DSIC a posteriori to the click realizations. Also, in each of the above cases, note that the utility of an agent i, $u_i(v_i, b_{-i}, \rho, t) \geq 0$. Therefore, by truthful bidding he never gets a negative utility. This proves that our mechanism is individually rational.

We next discuss the regret incurred by Δ-UCB. In order to prove the regret results, we will first need to prove several other lemmas.

Lemma 1. *Social Welfare UCB index: For an agent i, we define the social welfare UCB indices for agent i as,*

$$\widehat{W}^{+}_{i,t} = \widehat{\mu}_{i,t} v_i + \epsilon_{i,t} v_i = \widehat{\mu}_{i,t} v_i + \sqrt{2 \frac{v_i^2 \log T}{N_{i,t}}} \tag{3}$$

$$\widehat{W}^{-}_{i,t} = \widehat{\mu}_{i,t} v_i - \epsilon_{i,t} v_i = \widehat{\mu}_{i,t} v_i - \sqrt{2 \frac{v_i^2 \log T}{N_{i,t}}} \tag{4}$$

Then, $\forall t\ P\left(\left\{\omega : W_i \notin [\widehat{W}^{-}_{i,t}(\omega), \widehat{W}^{+}_{i,t}(\omega)]\right\}\right) \leq T^{-4}$.

Proof. Let $\widehat{\mu}^{+}_{i,t}$ and $\widehat{\mu}^{-}_{i,t}$ denote the UCB and LCB indices for the estimate $\widehat{\mu}_i$. Then the events $\{\omega : \mu_i \notin [\widehat{\mu}^{-}_{i,t}(\omega),\ \widehat{\mu}^{+}_{i,t}(\omega)]\}$, $\{\omega : W_i \notin [\widehat{W}^{-}_{i,t}(\omega), \widehat{W}^{+}_{i,t}(\omega)]\}$ are identical. So, $P(W_i \notin [\widehat{W}^{-}_{i,t}, \widehat{W}^{+}_{i,t}]) = P(\mu_i \notin [\widehat{\mu}^{-}_{i,t}, \widehat{\mu}^{+}_{i,t}])$. An application of Hoeffding bound [12] gives $P(\mu_i \notin [\widehat{\mu}^{-}_{i,t}, \widehat{\mu}^{+}_{i,t}]) \leq \exp(-2N_{i,t}\epsilon^2_{i,t})$. As per the mechanism $\epsilon_{i,t} = \sqrt{2 \log T / N_{i,t}}$. So $P(\mu_i \notin [\widehat{\mu}^{-}_{i,t}, \widehat{\mu}^{+}_{i,t}]) \leq \exp(-2N_{i,t} \times 2 \log T / N_{i,t}) = T^{-4}$.

Lemma 2. *Suppose at time step t, $N_{i,t} > \frac{8 v^2_{max} \log T}{\Delta^2}$ $\forall i \in [K]$. Then $\forall i \in [K]$, $2\epsilon_{i,t} v_i < \Delta$.*

Proof. Given that $N_{i,t} > \frac{8 v^2_{max} \log T}{\Delta^2}$. Therefore,

$$\Delta^2 > \frac{8 v^2_{max} \log T}{N_{i,t}} \geq \frac{8 v_i^2 \log T}{N_{i,t}} \geq 4 \left[\frac{2 v_i^2 \log T}{N_{i,t}}\right]$$

Taking square roots on both sides of the above equation yields $\Delta > 2\epsilon_{i,t} v_i$ thereby proving the lemma.

Lemma 3. *For an agent i and time step t, let $B_{i,t}$ be the event $B_{i,t} = \{\omega : W_i \notin [\widehat{W}^-_{i,t}, \widehat{W}^+_{i,t}]\}$. Define the event $G = \bigcap_t \bigcap_{i\in[K]} B^c_{i,t}$, where $B^c_{i,t}$ is the complement of $B_{i,t}$. Then $P(G) \geq 1 - \frac{1}{T^2}$.*

Proof. From Lemma 1, the probability of the 'bad' event, $P(B_{i,t}) \leq T^{-4}$.

$$\begin{aligned} P(G) &= P\left(\bigcap_t \bigcap_i B^c_{i,t}\right) = 1 - P\left(\left(\bigcap_t \bigcap_i B^c_{i,t}\right)^c\right) \\ &= 1 - P\left(\bigcup_t \bigcup_i B_{i,t}\right) = 1 - \sum_t \sum_{i\in[K]} P(B_{i,t}) \\ &\geq 1 - \sum_t \sum_{i\in[K]} T^{-4} \geq 1 - \frac{1}{T^2} \end{aligned}$$

The last statement follows by summing over all rounds and using the fact that $K \ll T$.

Theorem 2. *Suppose at time step t, $N_{j,t} > \frac{8v^2_{max}\log T}{\Delta^2}$ $\forall j \in [K]$. Then $\forall i \in [K] \setminus S_\Delta$, $\widehat{W}^+_{i_*,t} > \widehat{W}^+_{i,t}$ with high probability ($= 1 - 2/T^4$).*

Proof: In Theorem 1, we have shown that Δ-UCB is DSIC. Therefore, all the agents bid their valuations truthfully, $b_i = v_i$ $\forall i \in [K]$. Suppose in exploitation round t, a sub-optimal arm is pulled. Therefore, $\widehat{W}^+_{i,t} \geq \widehat{W}^+_{i_*,t}$. Then one of the following three conditions must have happened.
Condition 1: $W_i < \widehat{W}^-_{i,t}$. This condition implies a drastic overestimate of the sub-optimal arm i so that the true social welfare W_i is even below the LCB index $\widehat{W}^-_{i,t}$. The figure below shows this case.

Condition 2: $W_* > \widehat{W}^+_{i_*,t}$. This implies an underestimate of the optimal arm so that the true social welfare W_* lies above even the UCB index $\widehat{W}^+_{i_*,t}$. See subsequent figure.

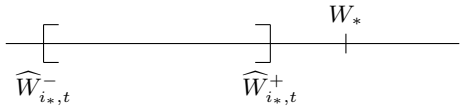

Condition 3: $W_* - W_i < 2\epsilon_{i,t} v_i$**.** This implies an overlap in the confidence intervals of the optimal and sub-optimal arm. Even though Conditions 1 and 2 are false, still the UCB of sub-optimal arm i is greater than the UCB of the optimal arm i_*.

From the above figure, $W_* - W_i \leq \widehat{W}^+_{i,t} - \widehat{W}^-_{i,t} \leq 2\epsilon_{i,t}$
If all the three conditions above were false, then,

$$\widehat{W}^+_{i_*,t} > W_* > W_i + 2\epsilon_{i,t} v_i > \widehat{W}^-_{i,t} + 2\epsilon_{i,t} v_i = \widehat{W}^+_{i,t}$$

This implies that $\widehat{W}^+_{i_*,t} > \widehat{W}^+_{i,t}$, leading to a contradiction.

As per the statement of the theorem, $N_{i,t} > \frac{8 v_{max}^2 \log T}{\Delta^2}$. Therefore by Lemma 2, $2\epsilon_{i,t} v_i < \Delta$. For $i \in [K] \setminus S_\Delta$, $W_* - W_i > \Delta > 2\epsilon_{i,t} v_i$. So Condition 3 above does not hold true. So if the sub-optimal arm i must have been pulled, only possibilities are for Condition 1 or 2.

$$\begin{aligned} P(\widehat{W}^+_{i,t} > \widehat{W}^+_{i_*,t}) &\leq P(\text{Condition 1}) + P(\text{Condition 2}) \\ &\leq P(B_{i,t}) + P(B_{i_*,t}) \leq 2/T^{-4} \end{aligned}$$

$$P(\widehat{W}^+_{i_*,t} > \widehat{W}^+_{i,t}) = 1 - P(\widehat{W}^+_{i,t} > \widehat{W}^+_{i_*,t}) \geq 1 - \frac{2}{T^4}$$

thereby completing the proof.

We are now ready to state our main result on the incurred regret.

Theorem 3. *If the Δ-UCB mechanism is executed for a total time horizon of T rounds, it achieves an expected Δ-regret of $O(\log T)$.*

Proof. The main idea in the proof is to compute the Δ-regret conditional on two events - G and G^c and then to find a bound for these two conditional expectations.

$$\begin{aligned} \mathbb{E}\left[\Delta\text{-regret}|G\right] &= \mathbb{E}\left[\Delta\text{-regret}|\forall t, \forall i\ W_i \in [\widehat{W}^-_{i,t}, \widehat{W}^+_{i,t}]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{T} (W_* - W_{I_t})\, \mathbb{1}\left[I_t \in [K] \setminus S_\Delta\right] |\forall t, \forall i\ W_i \in [\widehat{W}^-_{i,t}, \widehat{W}^+_{i,t}]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{T} (W_* - W_{I_t})\, \mathbb{1}\left[I_t \in [K] \setminus S_\Delta\right] |W_{I_t} \in [\widehat{W}^-_{I_t,t}, \widehat{W}^+_{I_t,t}]\right] \\ &= \frac{8K v_{max}^3 \log T}{\Delta^2} \end{aligned}$$

The last step comes from the fact that Conditions 1 and 2 in the proof of Theorem 2 are eliminated as we are given that the event G has occurred. After exploration rounds, $N_{i,t} \geq 8Kv_{max}^2 \log T/\Delta^2$. From Theorem 2, no Δ-regret occurs during exploitation since G is true. Therefore the regret is only incurred during the exploration rounds.

We now compute $\mathbb{E}[\Delta\text{-regret}|G^c]$.

$$\mathbb{E}[\Delta\text{-regret}|G^c] \leq Tv_{max} \tag{5}$$

But $P(G^c) = 1 - P(G) < \frac{1}{T^2}$ from Lemma 3.
Putting all the steps together,

$$\begin{aligned}
\mathbb{E}[\Delta\text{-regret}] &= \mathbb{E}[\Delta\text{-regret}|G]\,P(G) + \mathbb{E}[\Delta\text{-regret}|G^c]\,P(G^c) \\
&\leq \frac{8Kv_{max}^3 \log T}{\Delta^2} * 1 + Tv_{max} * \frac{1}{T^2} \\
&\leq \frac{8Kv_{max}^3 \log T}{\Delta^2} + 1
\end{aligned} \tag{6}$$

The second term is less than 1 as $v_{max} \ll T$. This completes the proof.

6 Extension to Multi-Slot SSA

In the previous sections, we assumed that there was a single slot for which the advertisers were competing. We now look at a more general setting where there are M slots to be allocated to the K agents. As before, each advertiser has exactly one ad for display and the CTR for advertisement i is denoted by μ_i. Recall that in the case of single slot auctions, the CTR exactly denoted the probability with which an ad received a click. However in the generalized setting of multi-slot auctions, an additional parameter comes into play while computing the click probability.

Each position or slot m is associated with a parameter λ_m called 'prominence'. λ_m denotes the probability with which a user observes an ad at slot $m+1$ given he has observed the ad at slot m. In order to understand the need for this parameter, a useful scenario to imagine is the listing of web-pages in Google for a query. There are two phases that one can think of once the listing of pages or results have appeared.

Phase 1: This is the phase where a user scans through the pages listed. A page listed higher up in the ranking (say second from the top) has more chances of being observed by a user rather than a page that is far below in the ranking (say fifth from the top). λ_4, for instance, denotes the probability that a user observes the fifth page, given he has observed the fourth page. Coming back to sponsored ads, we assume that $\lambda_1 = 1$, that is, the ad listed in the first slot is surely observed. We denote by Γ_m the probability that an ad at slot m is observed. Γ_m is computed as, $\Gamma_m = \prod_{s=1}^{m-1} \lambda_s$.

Phase 2: After having scanned through the list, the user decides to click one or more of the shown ads. In the multi-slot setting [11], it is assumed that multiple

ads in a listing may receive clicks. The probability that ad i receives a click when shown at slot $m = \Gamma_m \mu_i$.

We assume that λ_m, $m = 1, \ldots, M$ are known to the planner a-priori. The problem of learning these parameters along with the CTR μ is much harder due to the well-known identifiability problem. Approaches such as Expectation Maximization (EM) [6] may be used but the guarantees for such approaches are still open. Hence we work in the setting where the λs and hence Γs are known.

The above modeling assumptions are as per standard conventions [10]. In the multi-slot setting, the allocation is given to multiple agents at every time step. We denote by $\mathcal{A}(b, \rho, t) \subset \{1, \ldots, K\}$, the allocation at time t for bids b and click realization ρ. The cardinality of the allocated set $|\mathcal{A}(b, \rho, t)| = M$. We also use the notation $\mathcal{A}_i(b, \rho, t) = m$ to denote the allocation to agent i at time t is slot m, for the bid profile b, click realization ρ. If an agent i is not allocated any of the M slots at time t, we say $\mathcal{A}_i(b, \rho, t) = 0$.

We denote by $W_{i,m}$ the social welfare of agent i, when he is given slot m. $W_{i,m}$ is the expected valuation an agent i receives when he is given slot m and is computed as,

$$W_{i,m} = \Gamma_m \mu_i v_i \tag{7}$$

Symbol	Description
M	No. of slots
$[M]$	Set of M slots
λ_m	Prominence (Probability with which a user observes an ad at slot $m+1$ given he has observed the ad at slot m)
Γ_m	Probability that an ad at slot m is observed
$W_{i,m}$	Social welfare when agent i is allocated slot m
$N_{i,t}$	No. of times agent i has been selected till time t over all slots
$K^{(m)}$	Optimal agent for slot m
$W_{*,m}$	Social welfare when agent $K^{(m)}$ is allocated slot m
$S_{\Delta,m}$	Set of agents whose social welfare is less than Δ away from $K^{(m)}$. These agents do not contribute to Δ-regret when allocated slot m.

Table 2. Additional notations for multi-slot SSA

Having described the multi-slot setting, we now analyze the scenario from the view of the search engine or central planner. In the ideal scenario, the planner would like to allot the ads exactly to the top M agents with the largest social welfare. This use case has been studied in the literature [11] and exploration separated mechanisms with regret of $O(T^{2/3})$ have been proposed. Various possible allocations are explored for $O(T^{2/3})$ time steps for every agent after which the allocation algorithm is guaranteed to converge to the ideal allocation with high probability. As in the single slot case, $O(T^{2/3})$ exploration rounds are required

to distinguish all the agents perfectly from each other, when there are agents whose social welfare values are arbitrarily close.

However, a much more practical problem of interest is to study and design mechanisms when the search engine is indifferent to a gap in Δ in social welfare for every slot. We observe that in cases where the agents are well-separated, $O(T^{2/3})$ exploration rounds are not required. In fact, $O(\log T)$ exploration rounds are sufficient to converge to an allocation that is well within the requirements of the search engine.

Having explained the problem, we now formalize the notions of separatedness in this setting. Let $K^{(1)}, \ldots, K^{(M)} \in [K]$ be the best M agents in terms of their single slot social welfare values, that is, $\mu_{K^{(1)}} v_{K^{(1)}} > \mu_{K^{(2)}} v_{K^{(2)}} > \ldots > \mu_{K^{(M)}} v_{K^{(M)}}$. Let $W_{*,m} = W_{K^{(m)},m}$. The ideal solution would be to allocate agent $K^{(m)}$ the slot m. This allocation would yield the largest social welfare but in the worst case, when the agents' social welfares are separated by a function of T, converging to this optimal allocation would require $O(T^{2/3})$ exploration rounds [11]. Instead, for a prescribed value of Δ fixed by the search engine, define the set,

$$S_{\Delta,m} = \left\{ i \in [K] : W_{K^{(m)},m} - W_{i,m} < \Delta \right\}. \tag{8}$$

$S_{\Delta,m}$ is the set of all agents whose social welfare is at most Δ away from the agent $K^{(m)}$ (who should have ideally been given slot m). The planner is indifferent to the regret contributed by the agents in $S_{\Delta,m}$, if any of them are allotted slot m. Hence we define the multi-slot Δ-regret metric as,

$$\Delta\text{-regret} = \sum_{t=1}^{T} \sum_{m=1}^{M} (W_{*,m} - W_{I_{t,m},m}) \mathbb{1}\left[I_{t,m} \in [K] \setminus S_{\Delta,m} \right]$$

The Δ-UCB mechanism for the multi-slot SSA is given in Algorithm 2.

We analyze the regret and truthfulness of Algorithm 2. The lemmas and theorems for establishing the results for the multi-slot setting are similar to the single slot setting, however there are subtle differences in proving many of the results. We will highlight them as and when necessary.

Theorem 4. *In the multi-slot setting $\Delta - UCB$ is Dominant Strategy Incentive Compatible (DSIC) and Individually Rational (IR). (Proof along the lines of Theorem 1).*

Lemma 4. *For an agent i and slot m, the social welfare UCB indices for agent i,*

$$\begin{aligned} \widehat{W}^{+}_{i,m,t} &= \Gamma_m \widehat{\mu}_{i,t} v_i + \epsilon_{i,m,t} \\ &= \Gamma_m \widehat{\mu}_{i,t} v_i + \sqrt{2 \frac{v_i^2 \Gamma_m^2 \log T}{N_{i,t}}} \end{aligned} \tag{9}$$

$$\begin{aligned} \widehat{W}^{-}_{i,m,t} &= \Gamma_m \widehat{\mu}_{i,t} v_i - \epsilon_{i,m,t} \\ &= \Gamma_m \widehat{\mu}_{i,t} v_i - \sqrt{2 \frac{v_i^2 \Gamma_m^2 \log T}{N_{i,t}}} \end{aligned} \tag{10}$$

satisfy $P(W_{i,m} \notin [\widehat{W}^{-}_{i,m,t}, \widehat{W}^{+}_{i,m,t}])) \leq T^{-4}\ \forall t$

Proof. The proof idea is similar to Lemma 1.

A noteworthy feature of our estimates is the following. An allocation of an ad i in a slot m yields a sample for the computation of not only $\widehat{W}_{i,m,t}$, but also for $\widehat{W}_{i,m',t}$ for all slots $m' \in \{1, \ldots, M\}$. This is because Γ_m is known to the planner a-priori. Therefore note that, the number of allocations that ad i receives till time t, $N_{i,t}$ is the sum of the number of allocations that agent i receives irrespective of the slot or inclusive of all the slots.

Lemma 5. *Suppose at time step* t, $N_{j,t} > \frac{8v^2_{max} \log T}{\Delta^2}\ \forall j \in [K]$. *Then* $\forall i \in [K]$ *and* $\forall m \in [M]$, $2\epsilon_{i,m,t} < \Delta$.

Proof. The proof is similar to Lemma 2.

Lemma 6. *For an agent* i, *slot* m *and time* t, *let* $B_{i,m,t}$ *be the event* $B_{i,m,t} = \{\omega : W_{i,m} \notin [\widehat{W}^{-}_{i,m,t}(\omega), \widehat{W}^{+}_{i,m,t}(\omega)]\}$. *Define the event* $G = \bigcap_t \bigcap_i \bigcap_m B^c_{i,m,t}$. *Then* $P(G) \geq 1 - \frac{1}{T^2}$.

Proof: The proof has some subtle differences from Lemma 3 because in the multi-slot extension, the events $B_{i,m,t}$ are not independent across the slots.
Observation: If an element ω from the set of outcomes is such that $\omega \in B_{i,m,t}$, then $\omega \in B_{i,m',t}\ \forall m' \in [M]$. This is because, for any two slots m and m',

$$\begin{aligned} W_{i,m} \notin [\widehat{W}^{-}_{i,m,t}, \widehat{W}^{+}_{i,m,t}] &\iff \mu_i \notin [\widehat{\mu}^{-}_{i,t}, \widehat{\mu}^{+}_{i,t}] \\ &\iff W_{i,m'} \notin [\widehat{W}^{-}_{i,m',t}, \widehat{W}^{+}_{i,m',t}] \end{aligned}$$

Therefore $P(\bigcup_m B_{i,m,t}) = P(B_{i,1,t})$. From Lemma 4,
$P(\bigcup_m B_{i,m,t}) = P(B_{i,1,t}) \leq T^{-4}$. Hence,

$$\begin{aligned} P(G) &= 1 - P\left(\bigcup_t \bigcup_i \bigcup_m B_{i,m,t}\right) = 1 - P\left(\bigcup_t \bigcup_i B_{i,1,t}\right) \\ &\geq 1 - \frac{1}{T^2} \text{ (inequality from Lemma 3).} \end{aligned}$$

Theorem 5. *Suppose at time* t, $N_{j,t} > 8v^2_{max} \log T/\Delta^2$
$\forall j \in [K]$. *Then* $\forall m \in [M], \forall i \in [K] \setminus S_{\Delta,m}$, $\widehat{W}^{+}_{K^{(m)},m,t} > \widehat{W}^{+}_{i,m,t}$ *with high probability* $(= 1 - 2/T^4)$.

Proof: Suppose at time t where $N_{j,t} > 8v^2_{max} \log T/\Delta^2\ \forall j \in [K]$,, there exists some $m \in [M]$ such that $\widehat{W}^{+}_{K^{(m)},m,t} < \widehat{W}^{+}_{i,m,t}$. (Note that this statement does not arise from any assumptions on the allocation, for instance, that agent i is given slot m. This is the major difference from Theorem 2). But the relation between the true social welfare values of these agents is $W_{K^{(m)},m} > W_{i,m}$. Then one of the following three conditions must have occurred, like in proof of Theorem 2.

Condition 1: $W_{i,m} < \widehat{W}^{-}_{i,m,t}$. This condition implies a drastic overestimate of the sub-optimal arm i so that the true mean social welfare $W_{i,m}$ is even below the LCB index $\widehat{W}^{-}_{i,m,t}$. The figure below captures this condition.

Condition 2: $W_{K^{(m)},m} > \widehat{W}^{+}_{K^{(m)},m,t}$. This implies an underestimate of the optimal arm so that the true mean social welfare $W_{K^{(m)},m}$ lies above even the UCB index $\widehat{W}^{+}_{K^{(m)},m,t}$. See figure below.

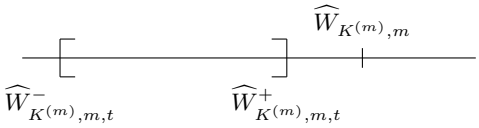

Condition 3: $W_{K^{(m)},m} - W_{i,m} < 2\epsilon_{i,m,t}$. This implies an overlap in the confidence intervals of the optimal and sub-optimal arm. Even if, Conditions 1 and 2 are false, still the UCB of sub-optimal arm i is greater than the UCB of the optimal arm i_*.

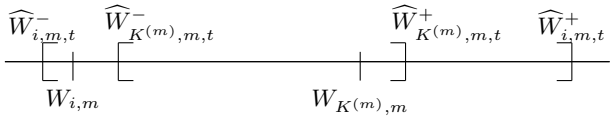

From the figure, $W_{K^{(m)},m} - W_{i,m} \leq \widehat{W}^{+}_{i,m,t} - \widehat{W}^{-}_{i,m,t} \leq 2\epsilon_{i,m,t}$. If all the three conditions above were false, then,

$$\begin{aligned}\widehat{W}^{+}_{K^{(m)},m,t} &> W_{K^{(m)},m} > W_{i,m} + 2\epsilon_{i,t} > \widehat{W}^{-}_{i,m,t} + 2\epsilon_{i,t} \\ &= \widehat{W}^{+}_{i,m,t} \quad (\text{ A contradiction!})\end{aligned}$$

As per the statement of the theorem, $N_{i,t} > 8v^2_{max} \log T/\Delta^2$. Therefore by Lemma 5, $2\epsilon_{i,m,t} < \Delta$. For agent $i \in [K] \setminus S_{\Delta,m}$, $W_{K^{(m)},m} - W_{i,m} > \Delta > 2\epsilon_{i,m,t}$. Therefore, Condition 3 above does not hold true. So,

$$\begin{aligned}P(\widehat{W}^{+}_{i,m,t} > \widehat{W}^{+}_{K^{(m)},m,t}) &\leq P(\text{Condition 1}) + P(\text{Condition 2}) \\ &\leq P(B_{i,m,t}) + P(B_{K^{(m)},m,t}) \leq 2/T^{-4}\end{aligned}$$

$$P(\widehat{W}^{+}_{K^{(m)},m,t} > \widehat{W}^{+}_{i,m,t}) = 1 - P(\widehat{W}^{+}_{i,m,t} > \widehat{W}^{+}_{K^{(m)},m,t})$$
$$\geq 1 - \frac{2}{T^4}$$

Theorem 6. *If the Δ-UCB mechanism is executed in the multiple slot scenario for a total time horizon of T rounds, it achieves an expected Δ-regret of $O(\log T)$.*

Proof. The proof idea has some subtle differences from the proof of Theorem 3. As before, we first compute the expected Δ-regret conditional on G. For the exploration rounds, the mechanism obtains a regret of $\xi = \frac{8MKv^3_{max}\log T}{\Delta^2}$.

$$\mathbb{E}\left[\Delta\text{-regret}|G\right] \leq \xi + \sum_{t=E}^{T}\sum_{m=1}^{M}(W_{K^{(m)},m} - W_{(I_{t,m}),m})\mathbb{1}\left[I_{t,m} \in K \setminus S_{\Delta,m}|G\right]$$

We will now show that the second term above evaluates to zero. For any m, the cardinality of $S_{\Delta,m}$ is at least m. This is because for all $K^{(j)}$ above m in the ranking of agents ($j < m$), $W_{K^{(m)},m} - W_{K^{(j)},m} < 0 < \Delta$ as $W_{K^{(j)},m} > W_{K^{(m)},m}$. Therefore there are at least $m-1$ agents in $S_{\Delta,m}$. Also $K^{(m)} \in S_{\Delta,m}$ as $W_{K^{(m)},m} - W_{K^{(m)},m} = 0 < \Delta$. Therefore $\forall j \in \{1,\ldots,m\}, K^{(j)} \in S_{\Delta,m}$. While allocating slot m, at least one of the agents in $S_{\Delta,m}$ must be free. This is by the pigeonhole principle. Now if the allocated agent for slot m, $I_{t,m} \in [K] \setminus S_{\Delta,m}$, one of the following two cases occur.

Case 1: The ideal agents $K^{(1)},\ldots,K^{(m-1)}$ for all the previous slots $1,\ldots,m-1$ have already been allocated before the allocation of slot m. This means that $K^{(m)}$ has not been allocated yet. Also, $\widehat{W}^{+}_{(I_{t,m}),m,\xi} > \widehat{W}^{+}_{K^{(m)},m,\xi}$. Since G is true and $t > \xi$, the above event cannot occur (by Theorem 5).

Case 2: The agent $K^{(m)}$ has already been allocated to some other slot before the allocation of slot m has begun. Therefore there is some agent $K^{(j)}, j < m$ with a larger social welfare value, who has still not been allocated. That is, $W_{K^{(j)},m} > W_{K^{(m)},m} > W_{(I_{t,m}),m}$. Given that $I_{t,m} \notin S_{\Delta,m}$. Therefore we can deduce that $I_{t,m} \notin S_{\Delta,j}$. This is because,

$$\begin{aligned}
&W_{K^{(m)},m} - W_{(I_{t,m}),m} \geq \Delta \\
&\implies W_{K^{(j)},m} - W_{(I_{t,m}),m} \geq \Delta \\
&\implies \mu_{K^{(j)}}v_{K^{(j)}} - \mu_{I_{t,m}}v_{I_{t,m}} \geq \Delta/\Gamma_m \\
&\implies \Gamma_j(\mu_{K^{(j)}}v_{K^{(j)}} - \mu_{I_{t,m}}v_{I_{t,m}}) \geq \Gamma_j\Delta/\Gamma_m \\
&\implies W_{K^{(j)},j} - W_{(I_{t,m}),j} \geq \Delta \qquad (11)
\end{aligned}$$

The last line in the above implications is true as $\Gamma_j > \Gamma_m$. But $\widehat{W}^{+}_{K^{(j)},m,\xi} < \widehat{W}^{+}_{(I_{t,m}),m,\xi}$. Then the inequality $\widehat{W}^{+}_{K^{(j)},j,\xi} < \widehat{W}^{+}_{(I_{t,m}),j,\xi}$ is also true due to the way the slot specific UCB indices are computed. From Theorem 5 for slot j, we find that $\widehat{W}^{+}_{K^{(j)},j,\xi} > \widehat{W}^{+}_{(I_{t,m}),j,\xi}$. Therefore we get that $\mathbb{E}\left[\Delta\text{-regret}|G\right] \leq \xi$.

Also, $P(G^c) = 1 - P(G) < \frac{1}{T^2}$ from Lemma 6.
Putting all the steps together,

$$\begin{aligned}\mathbb{E}\left[\Delta\text{-regret}\right] &= \mathbb{E}\left[\Delta\text{-regret}|G\right] P(G) \\ &\quad + \mathbb{E}\left[\Delta\text{-regret}|G^c\right] P(G^c) \\ &\leq \frac{8KMv_{max}^3 \log T}{\Delta^2} * 1 + TMv_{max} * \frac{1}{T^2} \\ &\leq \frac{8KMv_{max}^3 \log T}{\Delta^2} + v_{max} \end{aligned} \tag{12}$$

The simplification in the second line is because $\mathbb{E}\left[\Delta\text{-regret}|G^c\right] \leq TMv_{max}$. In the last line we use the fact that $M \ll T$. This completes the proof.

7 Conclusion

For the first time, we have addressed the realistic use case in MAB mechanisms where a planner has the option to specify a tolerance level for sub-optimal arms. The metric Δ-regret that we have introduced captures the notion of regret in this scenario. We have provided a deterministic, truthful, exploration separated MAB mechanism Δ-UCB which achieves a Δ-regret of $O(\log T)$. In particular we have analyzed the mechanism for sponsored search auctions for both single slot SSA and multi-slot SSA. The results are generic and will apply equally well to other applications as well.

References

1. S. Agrawal and N. Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *COLT*, pages 39.1–39.26, 2012.
2. P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
3. M. Babaioff, R. D. Kleinberg, and A. Slivkins. Truthful mechanisms with implicit payment computation. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC'10)*, pages 43–52. ACM, 2010.
4. M. Babaioff, Y. Sharma, and A. Slivkins. Characterizing truthful multi-armed bandit mechanisms. *SIAM Journal on Computing*, 43(1):194–230, 2014.
5. S. Bhat, D. Padmanabhan, S. Jain, and Y. Narahari. A truthful mechanism with biparameter learning for online crowdsourcing: (extended abstract). In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS'16), Singapore, May 9-13, 2016*, pages 1385–1386, 2016.
6. C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
7. A. Biswas, S. Jain, D. Mandal, and Y. Narahari. A truthful budget feasible multi-armed bandit mechanism for crowdsourcing time critical tasks. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS'15)*, pages 1101–1109, 2015.
8. S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.

9. N. R. Devanur and S. M. Kakade. The price of truthfulness for pay-per-click auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC'09)*, pages 99–106, 2009.
10. N. Gatti, A. Lazaric, M. Rocco, and F. Trovò. Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. *Artificial Intelligence*, 227:93–139, 2015.
11. N. Gatti, A. Lazaric, and F. Trovò. A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12)*, pages 605–622, 2012.
12. W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
13. S. Jain, S. Bhat, G. Ghalme, D. Padmanabhan, and Y. Narahari. Mechanisms with learning for stochastic multi-armed bandit problems. *Indian Journal of Pure and Applied Mathematics*, 47(2):229–272, 2016.
14. S. Jain, G. Ghalme, S. Bhat, S. Gujar, and Y. Narahari. A deterministic MAB mechanism for crowdsourcing with logarithmic regret and immediate payments. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS'16), Singapore, May 9-13, 2016*, pages 86–94, 2016.
15. R. B. Myerson. Game theory: analysis of conflict. *Harvard University*, 1991.
16. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

Algorithm 2 Δ-UCB Mechanism for multiple slot SSA

Input:

M : No. of slots, K: No. of agents, T: Time horizon

Δ : parameter fixed by the center

v_{max} : Maximum valuation of the agents

Elicit bids $b = (b_1, b_2, \ldots, b_K)$ from all the agents

Initialize $\widehat{\mu}_{i,0} = 0, N_{i,0} = 0 \; \forall i \in [K]$,

$u = 8Kv_{max}^2 \log T/\Delta^2$

for $t = 1, \ldots, u$ **do** ▷ Exploration rounds

  $\mathcal{A}(b, \rho, t) = \phi$

  for $m = 1, \ldots, M$ **do**

    $I_{t,m} = (((t-1) \mod K) + m - 1) \mod K + 1$

    $N_{I_{t,m},t} = N_{I_{t,m},t-1} + 1$

    $\mathcal{A}(b, \rho, t) = \mathcal{A}(b, \rho, t) \cup I_{t,m}$ ▷ Allocate $I_{t,m}$ slot m and observe $\rho_{I_{t,m}}(t)$.

    $\widehat{\mu}_{I_{t,m},t} = (\widehat{\mu}_{I_{t,m},t-1} N_{I_{t,m},t-1} + \rho_{I_{t,m}}(t))/N_{I_{t,m},t}$

    $\epsilon_{I_{t,m},t} = \sqrt{2 \log T / N_{I_{t,m},t}}$

    $\widehat{\mu}^{+}_{I_{t,m},t} = \widehat{\mu}_{I_{t,m},t} + \epsilon_{I_{t,m},t}$

    $\widehat{\mu}^{-}_{I_t,t} = \widehat{\mu}_{I_{t,m},t} - \epsilon_{I_{t,m},t}$

  end for

  $\widehat{\mu}^{+}_{i,t} = \widehat{\mu}^{+}_{i,t-1}, \widehat{\mu}^{-}_{i,t} = \widehat{\mu}^{-}_{i,t-1} \; \forall i \in [K] \setminus \mathcal{A}(b, \rho, t)$

  $P_i^t(b, \rho) = 0 \; \forall i \in [K]$ ▷ Free rounds

end for

$\widehat{K}^{(1)}, \widehat{K}^{(2)}, \ldots, \widehat{K}^{(M)}, \ldots, \widehat{K}^{(K)}$ = sorted list of agents in the decreasing order of $\widehat{\mu}^{+}_{i,u} b_i$

for $t = u+1, \ldots, T$ **do** ▷ Exploitation rounds

  $\mathcal{A}(b, \rho, t) = \phi$

  for $m = 1, \ldots, M$ **do**

    $I_{t,m} = \widehat{K}^{(m)}$

    $\mathcal{A}(b, \rho, t) = \mathcal{A}(b, \rho, t) \cup \widehat{K}^{(m)}$

    $P^t_{\widehat{K}^{(m)}}(b, \rho) = \sum_{l=m+1}^{M+1} (\Gamma_{l-1} - \Gamma_l) \, \widehat{\mu}^{+}_{K^{(l)},u} b_{K^{(l)}} \rho_{\widehat{K}^{(m)}}(t)$

  end for

  $P_i^t(b, \rho) = 0 \; \forall i \in [K] \setminus \mathcal{A}(b, \rho, t)$

  $\widehat{\mu}^{+}_{i,t} = \widehat{\mu}^{+}_{i,u}, \widehat{\mu}^{-}_{i,t} = \widehat{\mu}^{-}_{i,u} \; \forall i \in [K]$ ▷ No more learning

end for